

**IN THE CLAIMS:**

1 1. (Currently Amended) A method for a storage operating system implemented in a  
2 storage system to concurrently perform readahead operations for a plurality of different  
3 read streams established in one or more files, directories, vdisks or luns stored in the  
4 storage system, the method comprising:

5 allocating at least one readset data structure ("readset") for each of the one or  
6 more files, directories, vdisks or luns in which the plurality of different read streams is  
7 established, wherein the number of readsets allocated for each file, directory, vdisk or lun  
8 depends on the size of that file, directory, vdisk or lun;

9 receiving a client read request at the storage system, the client read request  
10 indicating client-requested data for the storage operating system to retrieve from a file,  
11 directory, vdisk or lun stored in the storage system;

12 determining whether the received client read request matches any of the plurality  
13 of readsets allocated for the file, directory, vdisk or lun containing the client-requested  
14 data;

15 performing readahead operations in accordance with a set of readahead metadata  
16 stored in an associated readset that is determined to match the received client read  
17 request, wherein the readahead metadata describes the associated readset; and

18 if the received client read request does not match any of the readsets allocated for  
19 the file, directory, vdisk or lun containing the client-requested data, then performing the  
20 steps:

21 identifying the received client read request as being the first read  
22 request in a new read stream;

23 generating a set of readahead metadata associated with the new  
24 read stream;

25 selecting for reuse one of the readsets allocated for the file,  
26 directory, vdisk or lun containing the client-requested data; and

27 storing the generated set of readahead metadata associated with the  
28 new read stream in the readset selected for reuse.

1 2. (Previously Presented) The method of claim 1, further comprising:  
2 generating a separate set of readahead metadata for each of the plurality of  
3 different read streams; and  
4 storing each generated set of readahead metadata in a different readset allocated  
5 for the file, directory, vdisk or lun in which the read stream associated with the generated  
6 set of readahead metadata is established.

1  
1 3. (Original) The method of claim 1, further comprising:  
2 initializing each allocated readset to store a predetermined set of values.

1  
1 4. (Cancelled)

1  
1 5. (Previously Presented) The method of claim 2, wherein the number of readsets  
2 allocated for a file, directory, vdisk or lun is dynamically increased as the size of that file,  
3 directory, vdisk or lun is increased.

1  
1 6. (Original) The method of claim 1, wherein a first readset is determined to match the  
2 received client read request if the first readset stores a set of readahead metadata  
3 associated with a read stream that is extended by the client-requested data.

1  
1 7. (Original) The method of claim 1, wherein a second readset is determined to match  
2 the received client read request when the client-requested data is located within a  
3 predetermined fuzzy range associated with the second readset.

1  
1 8. (Original) The method of claim 7, wherein the fuzzy range is derived based on a  
2 multiple of a number of client-requested data blocks specified in the received client read  
3 request.

1  
1 9. (Original) The method of claim 7, wherein the fuzzy range extends in both a forward  
2 direction and a backward direction in relation to a last data block retrieved in a read  
3 stream associated with the second readset.

1 10. (Original) The method of claim 1, wherein a third readset is determined to match the  
2 received client read request if the third readset is determined to be unused.

1

1 11. (Original) The method of claim 10, wherein the third readset is determined to be  
2 unused when a level value stored in the third readset equals a special indicator value.

1

1 12. (Original) The method of claim 1, wherein readahead operations are not performed if  
2 the storage operating system determines that the file, directory, vdisk or lun containing  
3 the client-requested data is accessed using a random access style.

1

1 13. (Original) The method of claim 12, wherein a DAFS cache hint included in the  
2 received client read request indicates that the file, directory, vdisk or lun containing the  
3 client-requested data is accessed using a random access style.

1

1 14. (Original) The method of claim 1, wherein readahead operations are not performed  
2 unless:

3 (i) a readset is determined to match the received client read request; and

4 (ii) the matching readset stores a set of readahead metadata associated  
5 with a read stream that is extended by the client-requested data past a  
6 predetermined data block or memory address.

1

1 15. (Cancelled)

1

1 16. (Currently Amended) The method of claim 1, wherein the readset selected for reuse  
2 stores a level value that is less than or equal to level values stored in each of the other  
3 readsets associated with the file, directory, vdisk or lun containing the client-requested  
4 data.

1

1 17. (Original) The method of claim 1, wherein the client read request received at the  
2 storage system is a file-based client read request.

1

1 18. (Original) The method of claim 1, wherein the client read request received at the  
2 storage system is a block-based client read request.

1  
1 19-28 (Cancelled)

1 29. (Currently Amended) A storage system that employs a storage operating system to  
2 concurrently perform readahead operations for a plurality of different read streams  
3 established in one or more files, directories, vdisks or luns stored in the storage system,  
4 the storage system comprising:

5 means for allocating at least one readset data structure ("readset") for each of the  
6 one or more files, directories, vdisks or luns in which the plurality of different read  
7 streams is established, wherein the number of readsets allocated for each file, directory,  
8 vdisk or lun depends on the size of that file, directory, vdisk or lun;

9 means for receiving a client read request at the storage system, the client read  
10 request indicating client-requested data for the storage operating system to retrieve from a  
11 file, directory, vdisk or lun stored in the storage system;

12 means for determining whether the received client read request matches any of the  
13 plurality of readsets" allocated for the file, directory, vdisk or lun containing the client-  
14 requested data;

15 means for performing readahead operations in accordance with a set of readahead  
16 metadata stored in an associated readset that is determined to match the received client  
17 read request, wherein the readahead metadata describes the associated readset; and

18 if the received client read request does not match any of the readsets allocated for  
19 the file, directory, vdisk or lun containing the client-requested data, then means for  
20 performing:

21 means for identifying the received client read request as being the  
22 first read request in a new read stream;

23 means for generating a set of readahead metadata associated with  
24 the new read stream;

25 means for selecting for reuse one of the readsets allocated for the  
26 file, directory, vdisk or lun containing the client-requested data; and

27 means for storing the generated set of readahead metadata  
28 associated with the new read stream in the readset selected for reuse.

1

1 30. (Currently Amended) A computer-readable media comprising instructions for  
2 execution in a processor for the practice of a method for a storage operating system  
3 implemented in a storage system to concurrently perform readahead operations for a  
4 plurality of different read streams established in one or more files, directories, vdisks or  
5 luns stored in the storage system, the method comprising:

6 allocating at least one readset data structure ("readset") for each of the one or  
7 more files, directories, vdisks or luns in which the plurality of different read streams is  
8 established, wherein the number of readsets allocated for each file, directory, vdisk or lun  
9 depends on the size of that file, directory, vdisk or lun;

10 receiving a client read request at the storage system, the client read request  
11 indicating client-requested data for the storage operating system to retrieve from a file,  
12 directory, vdisk or lun stored in the storage system;

13 determining whether the received client read request matches any of the plurality  
14 of readsets allocated for the file, directory, vdisk or lun containing the client-requested  
15 data;

16 performing readahead operations in accordance with a set of readahead metadata  
17 stored in an associated readset that is determined to match the received client read  
18 request, wherein the readahead metadata describes the associated readset; and

19 if the received client read request does not match any of the readsets allocated for  
20 the file, directory, vdisk or lun containing the client-requested data, then performing the  
21 steps:

22 identifying the received client read request as being the first read  
23 request in a new read stream;

24 generating a set of readahead metadata associated with the new  
25 read stream;

26 selecting for reuse one of the readsets allocated for the file,  
27 directory, vdisk or lun containing the client-requested data; and

28                   storing the generated set of readahead metadata associated with the  
29                   new read stream in the readset selected for reuse.

1

1   31. (Currently Amended) A method for a storage operating system implemented in a  
2   storage system to concurrently perform readahead operations for a plurality of different  
3   read streams established in one or more files stored in the storage system, comprising:  
4         allocating at least one read set data structure ("readset") for each of the one or  
5   more files, directories, vdisks or luns in which the plurality of different read streams is  
6   established wherein the number of readsets allocated for each file depends on the size of  
7   that file;  
8         generating a separate set of readahead metadata for each of the plurality of  
9   different read streams; and  
10        storing each generated set of readahead metadata in a different readset allocated  
11   for the file in which the read stream associated with the generated set of readahead  
12   metadata is established;  
13        receiving a client read request at the storage system, the client read request  
14   indicating client-requested data for the storage operating system to retrieve from a file,  
15   stored in the storage system;  
16        determining whether the received client read request matches any of a plurality of  
17   readsets allocated for the file containing the client-requested data; and  
18        performing readahead operations in accordance with a set of readahead metadata  
19   stored in a readset that is determined to match the received client read request; and  
20        if the received client read request does not match any of the readsets allocated for  
21   the file, directory, vdisk or lun containing the client-requested data, then performing the  
22   steps:  
23                identifying the received client read request as being the first read  
24                request in a new read stream;  
25                generating a set of readahead metadata associated with the new  
26                read stream;  
27                selecting for reuse one of the readsets allocated for the file,  
28                directory, vdisk or lun containing the client-requested data; and

29 storing the generated set of readahead metadata associated with the  
30 new read stream in the readset selected for reuse.

1

1 32. (Previously Presented) The method of claim 31, wherein the file is broad term  
2 describing either a file, directory, vdisk or lun.

1

1 33. (Previously Presented) The method of claim 31, further comprising:  
2 initializing each allocated readset to store a predetermined set of values.

1

1 34. (Previously Presented) The method of claim 31, wherein the number of readsets  
2 allocated for a file is dynamically increased as the size of that file is increased.

1

1 35. (Previously Presented) The method of claim 31, wherein a first readset is determined  
2 to match the received client read request if the first readset stores a set of readahead  
3 metadata associated with a read stream that is extended by the client-requested data.

1

1 36. (Previously Presented) The method of claim 31, wherein a second readset is  
2 determined to match the received client read request when the client-requested data is  
3 located within a predetermined fuzzy range associated with the second readset.

1

1 37. (Previously Presented) The method of claim 36, wherein the fuzzy range is derived  
2 based on a multiple of a number of client-requested data blocks specified in the received  
3 client read request.

1

1 38. (Previously Presented) The method of claim 36, wherein the fuzzy range extends in  
2 both a forward direction and a backward direction in relation to a last data block retrieved  
3 in a read stream associated with the second readset.

1

1 39. (Previously Presented) The method of claim 31, wherein a third readset is determined  
2 to match the received client read request if the third readset is determined to be unused.

1

1 40. (Previously Presented) The method of claim 39, wherein the third readset is  
2 determined to be unused when a level value stored in the third readset equals a special  
3 indicator value.

1  
1 41. (Previously Presented) The method of claim 31, wherein readahead operations are not  
2 performed if the storage operating system determines that the file, directory, vdisk or lun  
3 containing the client-requested data is accessed using a random access style.

1  
1 42. (Previously Presented) The method of claim 41, wherein a DAFS cache hint included  
2 in the received client read request indicates that the file, directory, vdisk or lun containing  
3 the client-requested data is accessed using a random access style.

1  
1 43. (Previously Presented) The method of claim 31, wherein readahead operations are not  
2 performed unless:

- 3 (i) a readset is determined to match the received client read request; and  
4 (ii) the matching readset stores a set of readahead metadata associated  
5 with a read stream that is extended by the client-requested data past a  
6 predetermined data block or memory address.

1  
1 44. (Cancelled)

1  
1 45. (Currently Amended) The method of claim 31, wherein the readset selected for reuse  
2 stores a level value that is less than or equal to level values stored in each of the other  
3 readsets associated with the file, directory, vdisk or lun containing the client-requested  
4 data.

1  
1 46. (Previously Presented) The method of claim 31, wherein the client read request  
2 received at the storage system is a file-based client read request.

1  
1 47. (Previously Presented) The method of claim 31, wherein the client read request  
2 received at the storage system is a block-based client read request.